Gradient-Based Discrete Sampling: Algorithms and Applications

Ruqi Zhang

Assistant Professor

Purdue CS

Monte Carlo Seminar

October 14, 2025

Discrete Data and Models

- Discrete data

Text

- beginning in december 1934 , training exercises were conducted for the tetrarchs and their crews using hamilcar gliders
- beginning in march 1946 , training exercises were conducted by the tetrarchs and their crews with hamilcar gliders .
- beginning in may 1926 , training exercises were conducted between the tetrarchs and their crews using hamilcar gliders .
- beginning in late 1942 , training exercises were conducted with the tetrarchs and their crews onboard hamilcar gliders .
- beginning in september 1961 , training exercises were conducted between the tetrarchs and their crews in hamilcar gliders .

DNA

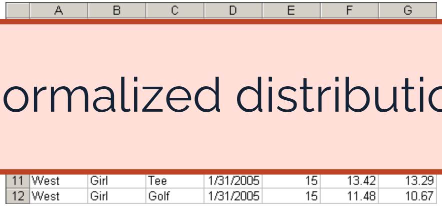

	A	B	C	D	E	F	G
11	West	Girl	Tee	1/31/2005	15	13.42	13.29
12	West	Girl	Golf	1/31/2005	15	11.48	10.67

Task: sample from an unnormalized distribution $\pi(\theta) \propto \exp(U(\theta))$

- Discrete models

Binary neural networks

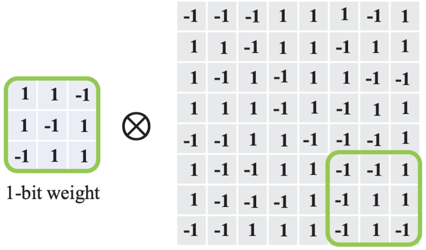

[Qin et al. 2020]

Discrete Samplers

- Gibbs sampling

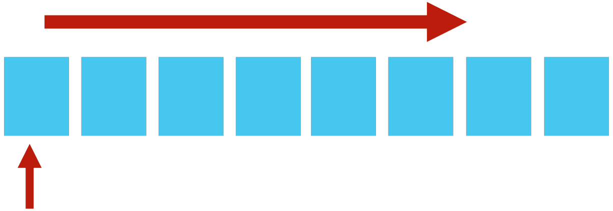

- Gibbs with Gradients

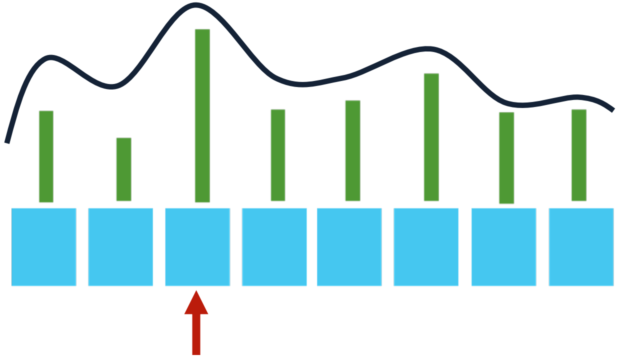

Oops I Took A Gradient: Scalable Sampling for Discrete Distributions. Grathwohl et al., ICML 2021

How to obtain gradient in discrete domains?

- Many common discrete unnormalized log-probability are differentiable functions

Distribution	$\log p(x) + \log Z$
Categorical	$x^T \theta$
Poisson[1]	$x \log \lambda - \log \Gamma(x+1)$
HMM	$\sum_{t=1}^{T} x_{t+1}^T A x_t - \frac{(w^T x - y)^2}{2\sigma^2}$
RBM	$\sum_i \text{softplus}(Wx+b)_i + c^T x$
Ising	$x^T W x + b^T x$
Potts	$\sum_{i=1}^{L} h_i^T x_i + \sum_{i,j=1}^{L} x_i^T J_{ij} x_j$
Deep EBM	$f_\theta(x)$

Oops I Took A Gradient: Scalable Sampling for Discrete Distributions. Grathwohl et al., ICML 2021

Discrete Samplers

- Gibbs sampling

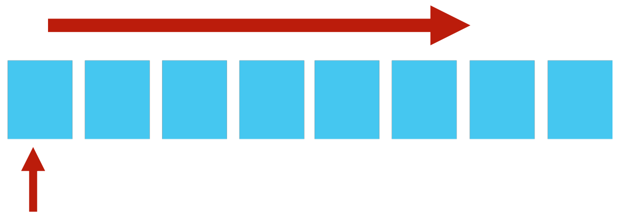

- Gibbs with Gradients

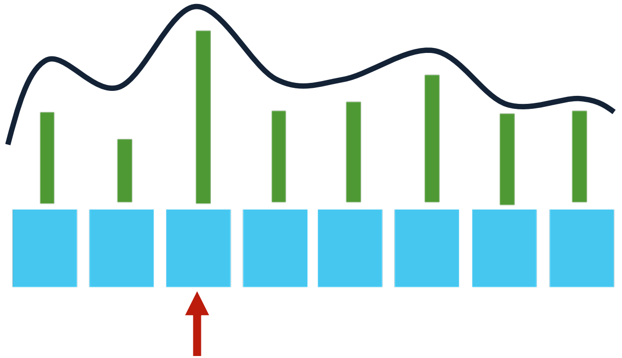

Only update one dim: suffer from high-dimensional and highly correlated distributions!

Oops I Took A Gradient: Scalable Sampling for Discrete Distributions. Grathwohl et al., ICML 2021

Continuous Sampler: Langevin Dynamics

$$\theta' = \theta + \frac{\alpha}{2}\nabla U(\theta) + \sqrt{\alpha}\xi, \qquad \xi \sim \mathcal{N}(0, I)$$

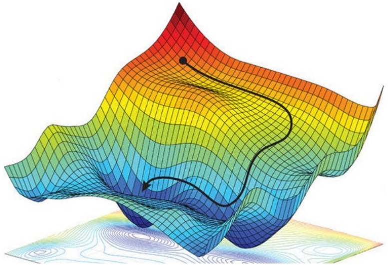

- Gradients guide the sampler to efficiently explore high probability regions
- Cheaply update all coordinates in parallel in a single step

What is the analog of Langevin dynamics in discrete domains?

Our Method: Discrete Langevin Proposal

$$q(\theta'|\theta) = \frac{\exp\left(-\frac{1}{2\alpha}\left\|\theta' - \theta - \frac{\alpha}{2}\nabla U(\theta)\right\|_2^2\right)}{Z_\Theta(\theta)}$$

- Langevin proposal is applicable to any kind of spaces
  - When $\Theta = \mathbb{R}^d$, recover the Gaussian proposal
  - When Θ is a discrete domain, obtain a gradient-based discrete proposal
- Coordinatewise factorization $q(\theta'|\theta) = \prod_{i=1}^{d} q_i(\theta'_i|\theta)$

$$\boxed{q_i(\theta'_i|\theta) = \text{Categorical}\left(\text{Softmax}\left(\frac{1}{2}\nabla U(\theta)_i(\theta'_i - \theta_i) - \frac{(\theta'_i - \theta_i)^2}{2\alpha}\right)\right)}$$

cheaply computed in parallel *Discrete Langevin Proposal* (DLP)

Visualization of Discrete Langevin Proposal

$$q_i(\theta_i'|\theta) = \text{Categorical}\Big(\text{Softmax}\Big(\frac{1}{2}\nabla U(\theta)_i(\theta_i' - \theta_i) - \frac{(\theta_i' - \theta_i)^2}{2\alpha}\Big)\Big)$$

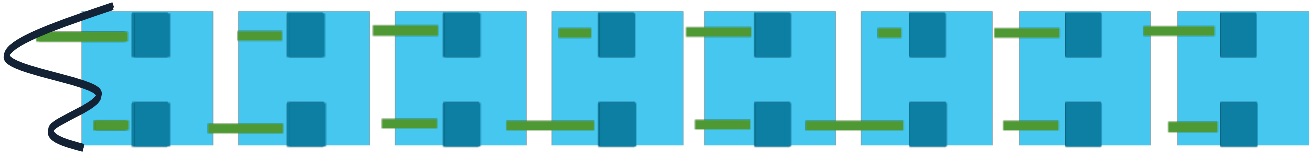

update all coordinates based on gradient info in parallel

Samplers: *discrete unadjusted Langevin algorithm* (DULA)

discrete Metropolis-adjusted Langevin algorithm (DMALA)

Visual Comparison

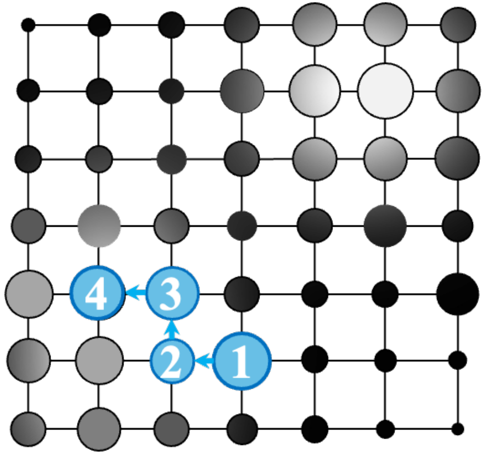

Existing discrete sampler

- Random walk
- Small move

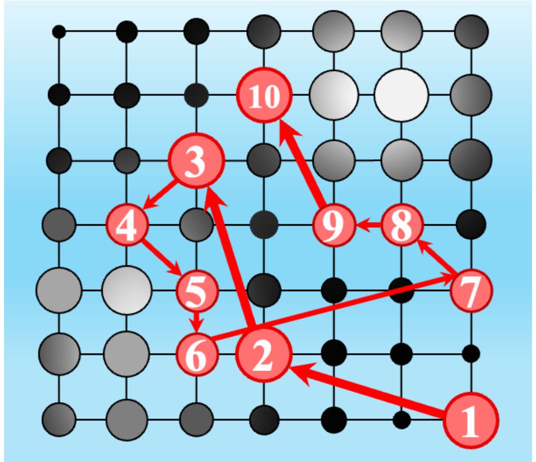

Discrete Langevin

- Gradient-informed exploration
- Large move

A Langevin-like Sampler for Discrete Distributions. R Zhang, X Liu, Q Liu. ICML 2022

Convergence Analysis

Theorem (informal): *The asymptotic bias of DULA's stationary distribution is zero for log-quadratic distributions and is small for distributions that are close to being log-quadratic*

Experiments: Restricted Boltzmann Machines

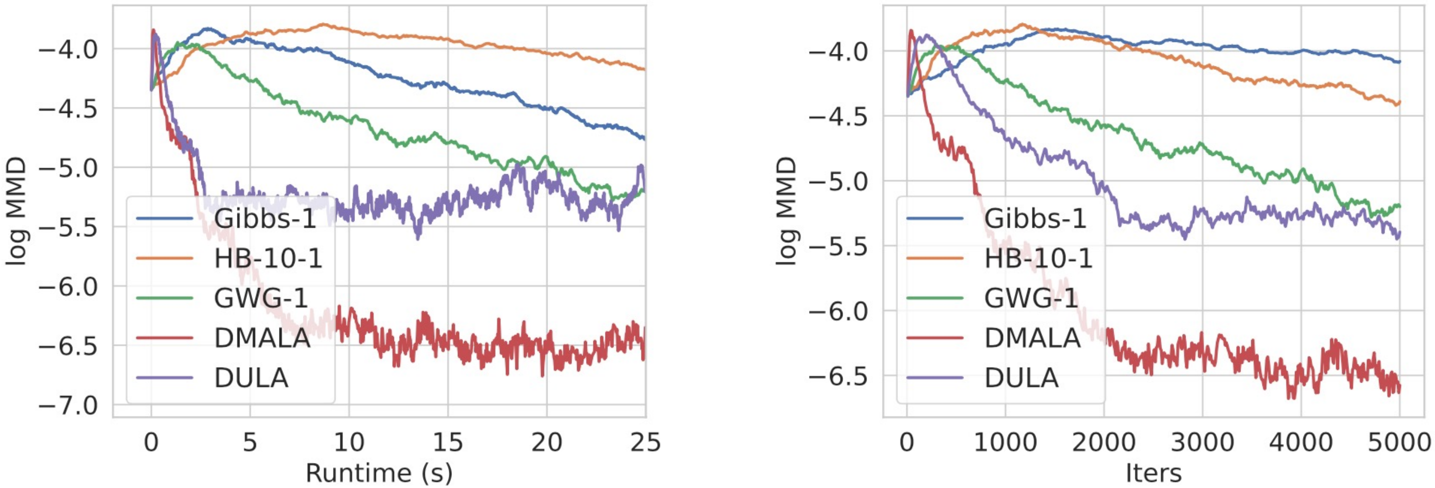

- DULA and DMALA converge faster to the target distribution

Experiments: Deep Energy-based Models

Dataset	VAE (Conv)	EBM (Gibbs)	EBM (GWG)	EBM (DULA)	EBM (DMALA)
Static MNIST	-82.41	-117.17	-80.01	-80.71	**-79.46**
Dynamic MNIST	-80.40	-121.19	-80.51	-81.29	**-79.54**
Omniglot	-97.65	-142.06	-94.72	-145.68	**-91.11**
Caltech Silhouettes	-106.35	-163.50	-96.20	-100.52	**-87.82**

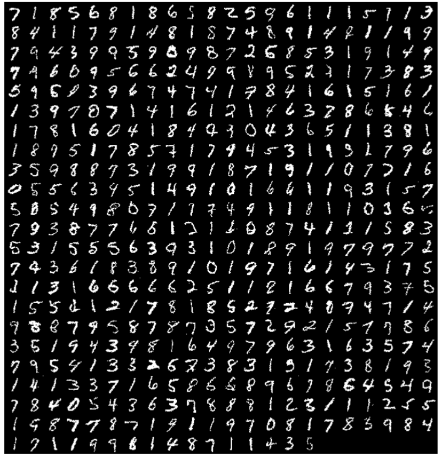
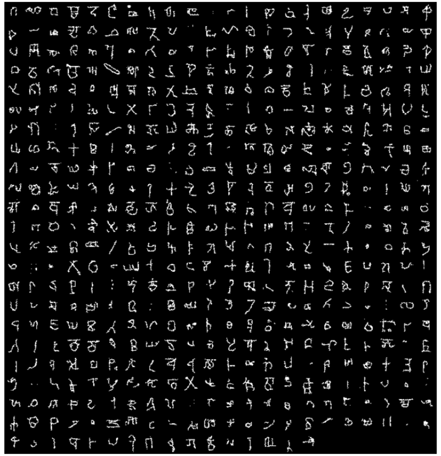
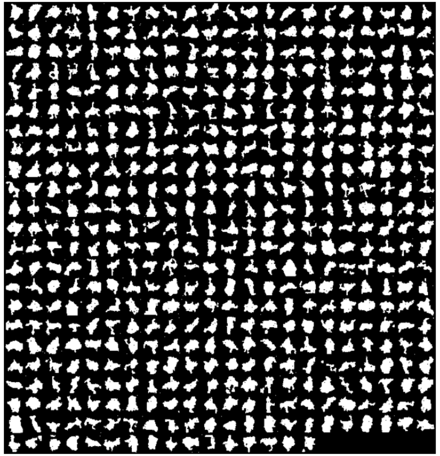

Image generation

Experiments: Language Models

Infilling Task: he had not , after all , [MASK] me the chance but [MASK] abandoned me [MASK] .

Gibbs Results:	**GWG Results**:	**DMALA Results**:
given me the chance but had abandoned me instead	given me the chance but had abandoned me instead	shown me the chance but had abandoned me anyway
given me the chance but had abandoned me instead	given me the chance but had abandoned me himself	shown me the chance but not abandoned me immediately
given me the chance but had abandoned me instead	offered me the chance but had abandoned me completely	gives me the chance but also abandoned me perhaps
given me the chance but had abandoned me completely	gave me the chance but had abandoned me anyway	grants me the chance but really abandoned me entirely
given me the chance but had abandoned me anyway	given me the chance but he abandoned me instead	offered me the chance but yet abandoned me instead

Model	Methods	Self-BLEU (↓)	Unique n-grams (%) (↑)						Corpus BLEU (↑)
			Self		WT103		TBC		
			$n=2$	$n=3$	$n=2$	$n=3$	$n=2$	$n=3$	
Bert-Base	Gibbs	86.84	10.98	16.08	18.57	32.21	21.22	33.05	**23.82**
	GWG	81.97	15.12	21.79	22.76	37.59	24.72	37.98	22.84
	DULA	**72.37**	**23.33**	**32.88**	27.74	**45.85**	30.02	**46.75**	21.82
	DMALA	72.59	23.26	32.64	**27.99**	45.77	**30.32**	46.49	21.85
Bert-Large	Gibbs	88.78	9.31	13.74	17.78	30.50	20.48	31.23	22.57
	GWG	86.50	11.03	16.13	19.25	33.20	21.42	33.54	**23.08**
	DULA	77.96	17.97	26.64	23.69	41.30	26.18	42.14	21.28
	DMALA	**76.27**	**19.83**	**28.48**	**25.38**	**42.94**	**27.87**	**43.77**	21.73

Text infilling

Extension: Multimodal Distributions

Cyclical stepsize and balancing parameter schedules

$$q_i(\theta'_i|\theta) = \mathrm{Cat}\left(\mathrm{Softmax}\left(\beta \nabla U(\theta)_i(\theta'_i - \theta_i) - \frac{(\theta'_i - \theta_i)^2}{2\alpha}\right)\right)$$

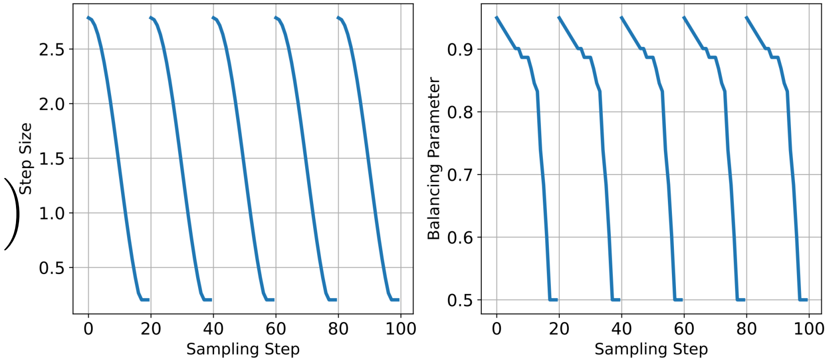

Theory: non-asymptotic convergence analysis based on TV distance

(a) Ground Truth (b) Random Walk (c) DMALA (d) AB (e) ACS (Ours)

Gradient-based Discrete Sampling with Automatic Cyclical Scheduling
P Pynadath, R Bhattacharya, A Hariharan, R Zhang. NeurIPS 2024

Extension: Combinatorial Optimization

Reheat mechanism:

- Detect when to reheat
- Increase the temperature to a predefined high value

Method	Type	SATLIB		ER-[700-800]		ER-[9000-11000]	
		Size ↑	Drop ↓	Size ↑	Drop ↓	Size ↑	Drop ↓
KaMIS	OR	425.96*	-	44.87*	-	381.31*	-
Gurobi	OR	425.95	0.00%	41.38	7.78%	N/A	N/A
Intel (Li et al., 2018a)	SL+TS	N/A	N/A	38.8	13.43%	N/A	N/A
	SL+G	420.66	1.48%	34.86	22.31%	284.63	25.35%
DGL (Böther et al., 2022)	SL+TS	N/A	N/A	37.26	16.96%	N/A	N/A
LwD(Ahn et al., 2020)	RL+S	422.22	0.88%	41.17	8.25%	345.88	9.29%
DIMES(Qiu et al., 2022)	RL+G	421.24	1.11%	38.24	14.78%	320.50	15.95%
	RL+S	423.28	0.63%	42.06	6.26%	332.80	12.72%
iSCO (Sun et al., 2023b)	S-1	422.65	0.78%	43.37	3.3%	377.44	1.0%
	S-32	424.16	0.42%	45.16	-0.6%	383.50	-0.5%
ReSCO(Ours)	S-1	422.76	0.75%	44.18	1.5%	378.25	0.8%
	S-32	**424.21**	**0.42%**	**45.24**	**-0.8%**	**383.75**	**-0.6%**

Reheated Gradient-based Discrete Sampling for Combinatorial Optimization. M Li, R Zhang. TMLR 2025

Application: LLM Control Generation

- Problem: struggle to balance fluency with constraint satisfaction

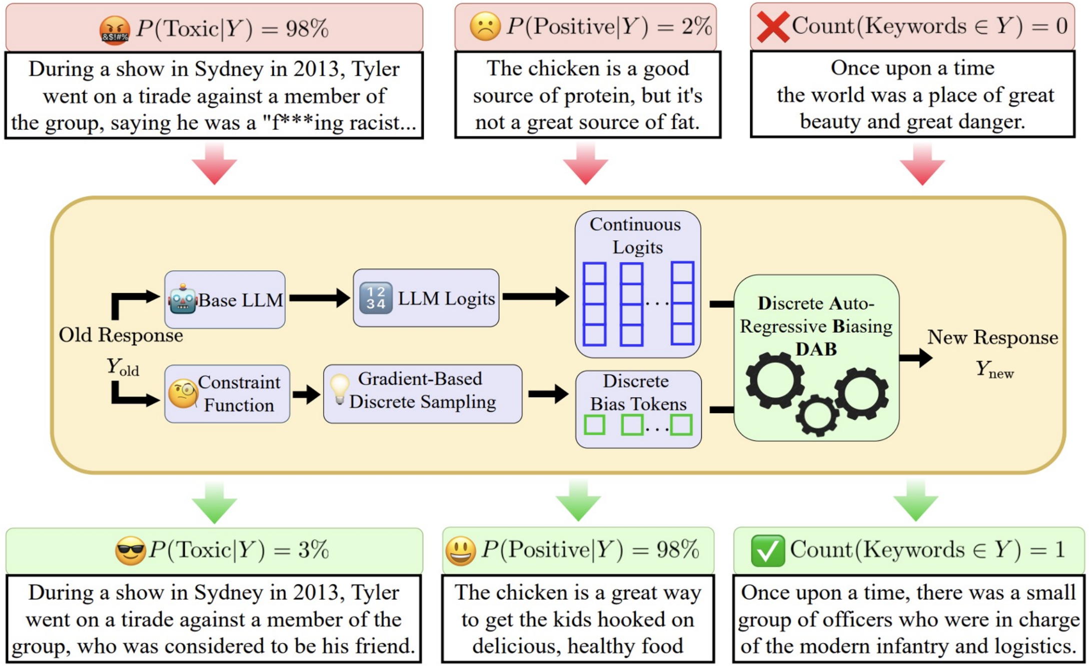

Controlled LLM Decoding via Discrete Auto-regressive Biasing. P Pynadath, R Zhang. ICLR 2025

Discrete Auto-regressive Biasing (DAB)

- Our joint target distribution:

$$P(Y, B|X) \propto P^{LM}(Y|X, B)\exp(f(B|X))$$

- X: query
- Y: response
- f: constraint function
- B: bias vectors

- How to sample?
  - Discrete Langevin within Gibbs

Controlled LLM Decoding via Discrete Auto-regressive Biasing. P Pynadath, R Zhang. ICLR 2025

Discrete Auto-regressive Biasing (DAB)

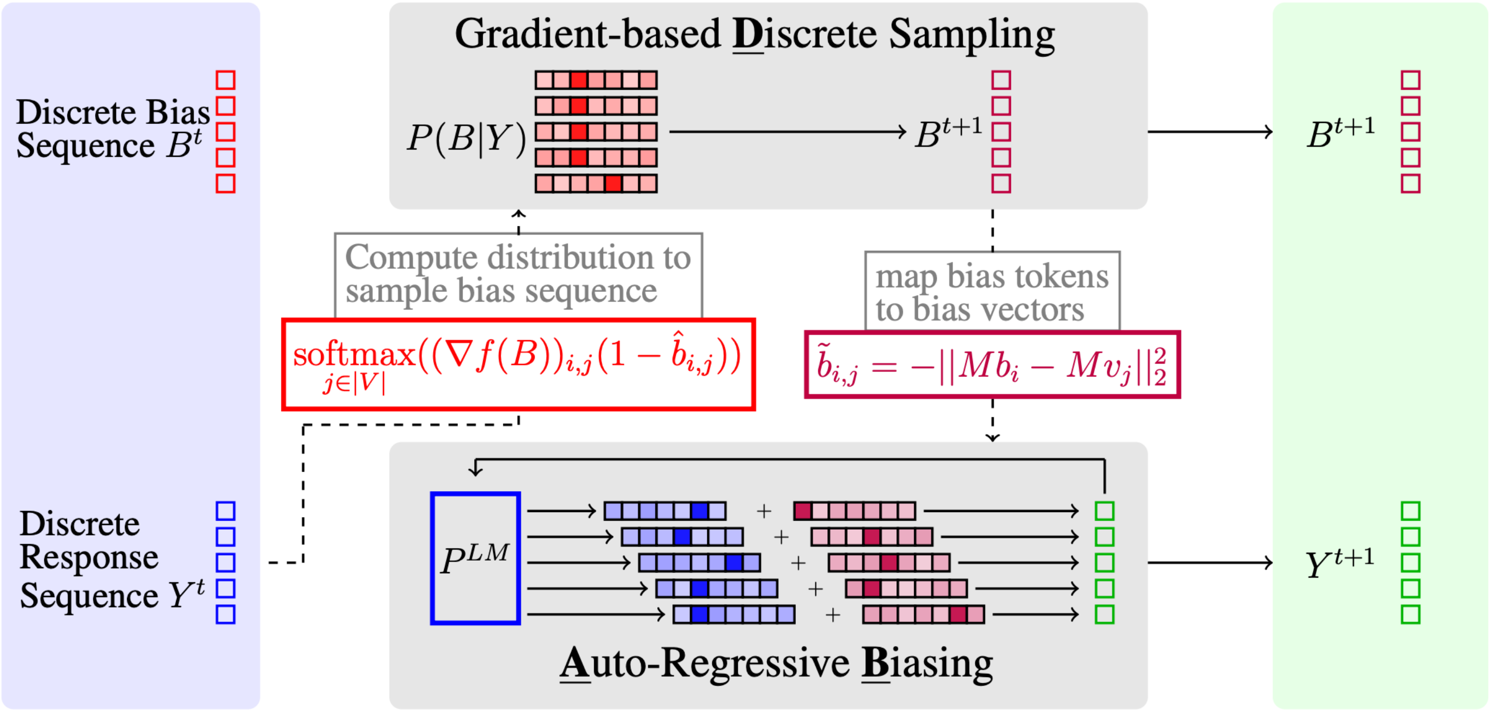

Controlled LLM Decoding via Discrete Auto-regressive Biasing. P Pynadath, R Zhang. ICLR 2025

DAB Results

		Control			Fluency	
Sentiment	*Int. Clsf* ↑	*Ext. Clsf (Yelp)* ↑	*Ext. Clsf (SST-2)* ↑	*CoLA* ↑	*REP-3gram* ↓	*PPL* ↓
MuCOLA	$.841 \pm .009$	<u>$.843 \pm .011$</u>	$.899 \pm .008$	$.681 \pm .008$	$.091 \pm .006$	34.786 ± 2.205
COLD	$.697 \pm .011$	$.515 \pm .015$	$.670 \pm .013$	$.731 \pm .008$	$.061 \pm .003$	$15.908 \pm .394$
BOLT	<u>$.903 \pm .006$</u>	$.747 \pm .013$	$.878 \pm .001$	$\mathbf{.874 \pm .005}$	$\mathbf{.0008 \pm .0002}$	$\mathbf{9.919 \pm .142}$
LM-Steer	-	$\mathbf{.900 \pm .008}$	<u>$.948 \pm .006$</u>	$.564 \pm .008$	$.117 \pm .007$	72.153 ± 3.195
DAB *(Ours)*	$\mathbf{.992 \pm .001}$	$\mathbf{.894 \pm .009}$	$\mathbf{.975 \pm .003}$	<u>$.860 \pm .005$</u>	<u>$.004 \pm .001$</u>	<u>$11.773 \pm .203$</u>
Toxicity	*Int. Clsf* ↓	*Avg. Max Toxicity* ↓	*Toxicity Pred. Prob.* ↓	*CoLA* ↑	*REP-3gram* ↓	*PPL* ↓
MuCOLA	$.098 \pm .002$	$.269 \pm .006$	7.6%	$.691 \pm .002$	$.006 \pm .001$	$58.015 \pm .435$
COLD	$.136 \pm .002$	$.266 \pm .007$	10.2%	$.667 \pm .001$	$.024 \pm .001$	$38.891 \pm .177$
BOLT	<u>$.065 \pm .001$</u>	<u>$.264 \pm .006$</u>	**6.8%**	$\mathbf{.830 \pm .001}$	$\mathbf{.001 \pm .0001}$	<u>27.283 ± 2.233</u>
LM-Steer	-	<u>$.265 \pm .006$</u>	<u>7.9%</u>	$.722 \pm .002$	<u>$.006 \pm .002$</u>	$52.697 \pm .356$
DAB *(Ours)*	$\mathbf{.057 \pm .001}$	$\mathbf{.211 \pm .006}$	**6.8%**	<u>$.806 \pm .001$</u>	$\mathbf{.001 \pm .0001}$	$\mathbf{25.609 \pm .126}$
Keyword	*BertScore* ↑	*Success Rate* ↑	-	*CoLA* ↑	*REP-3gram* ↓	*PPL* ↓
MuCOLA	$.8083 \pm .0004$	**100%**	-	$.248 \pm .004$	$.007 \pm .001$	475.301 ± 30.445
COLD	$.8123 \pm .0005$	**100%**	-	$.205 \pm .003$	$.020 \pm .001$	241.980 ± 4.943
BOLT	<u>$.8291 \pm .0003$</u>	99.1%	-	<u>$.705 \pm .006$</u>	<u>$.005 \pm .005$</u>	<u>32.019 ± 1.593</u>
DAB *(Ours)*	$\mathbf{.8303 \pm .0003}$	99.0%	-	$\mathbf{.726 \pm .005}$	$\mathbf{.004 \pm .001}$	$\mathbf{23.424 \pm .317}$

- Better fluency and constraint satisfaction trade-off
- 2x faster decoding time

Application: Molecular optimization & Drug design

- Discrete Langevin + Genetic algorithm

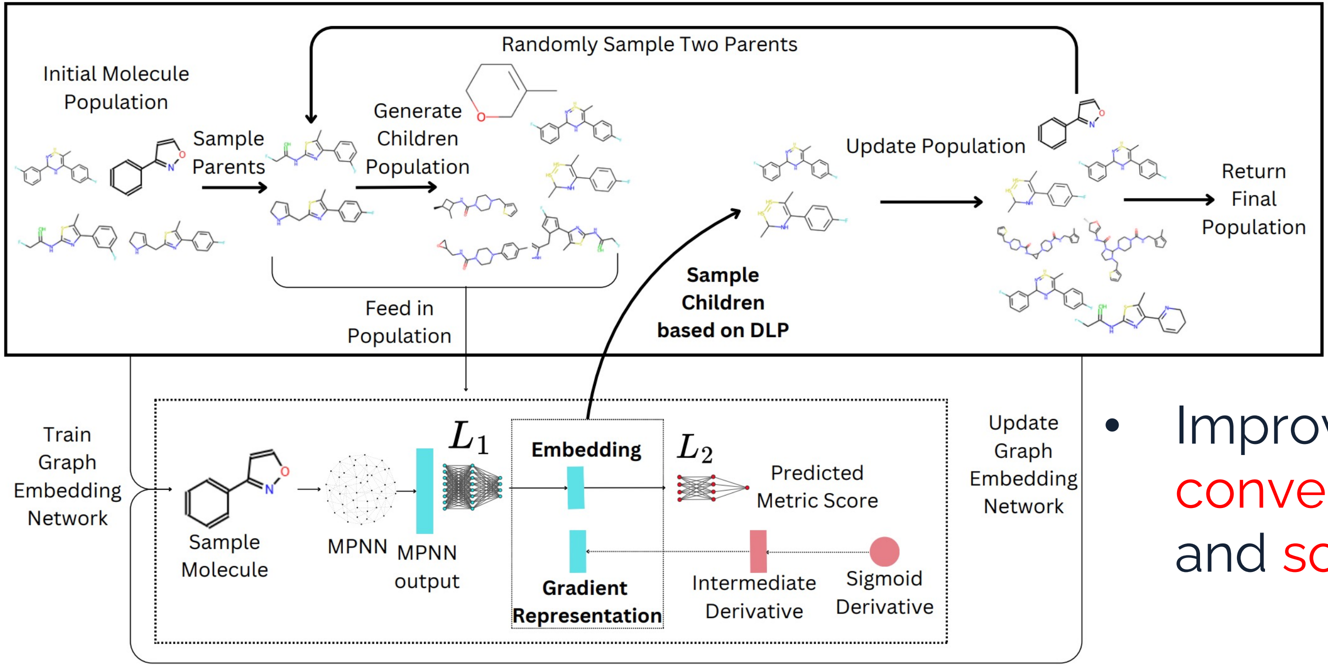

- Improve convergence speed and solution quality

Gradient GA: Gradient Genetic Algorithm for Drug Molecular Design
D Mukherjee, C Zhuang, Y Lu, T Fu, R Zhang. arXiv 2025

More Work on Gradient-based Discrete Sampling

- Without natural continuous extension

  Efficient Informed Proposals for Discrete Distributions via Newton's Series Approximation
  Y Xiang, D Zhu, B Lei, D Xu, R Zhang, AISTATS 2023

- Benchmark for discrete sampling: 7 samplers and 3 types of tasks

  DISCS: A Benchmark for Discrete Sampling
  K Goshvadi, H Sun, X Liu, A Nova, R Zhang, W Grathwohl, D Schuurmans, H Dai, NeurIPS 2023

- Diffusion language models

  Coming soon!

Takeaways

- Sampling in discrete domains can be very efficient by using a discrete version of Langevin dynamics
- Algorithm extensions:
  - Cyclical schedules for multimodal distributions
  - Reheat mechanism for combinatorial optimization
- Applications:
  - Classic models: Ising, Potts, Restricted Boltzmann Machines, Energy-based models
  - LLM control generation
  - Molecular optimization & Drug design

Thank you!